

IsoWAN: A NASA Science and Engineering Information and Services Framework

David J. Korsmeyer
dkorsmeyer@arc.nasa.gov
 NASA Ames Research Center

Edward T. Chow
edward.chow@jpl.nasa.gov
 NASA Jet Propulsion Laboratory

Michael P. Conroy
michael.conroy-1@ksc.nasa.gov
 NASA Kennedy Space Center

Abstract

We believe that the next evolutionary step in supporting wide-area application and services delivery to customers is a network framework that provides for collocation of applications and services at distinct sites in the network, an interconnection between these sites that is performance optimized for these applications, and value-added services for applications. We use the term IsoWAN to describe an advanced, isolated network interconnect services framework that will enable applications to be more secure, and able to access and be in use in both local and remote environments. The main functions of an IsoWAN are virtual localization of application services, an application service interface, coordinated delivery of applications and associated data to the customer, and supporting collaborative application development for customers. An initial pilot network between three NASA Centers; Ames Research Center, the Jet Propulsion Laboratory, and Marshall Space Flight Center, has been built and its properties will be discussed.

1: Introduction

“The ISE [Intelligent Synthesis Environment] aims to link scientists, design teams, manufacturers, suppliers, and consultants in the creation and operation of an aerospace system and in synthesizing its missions. The ultimate goal is to significantly increase creativity and knowledge and eventually dissolve rigid cultural boundaries among diverse engineering and science teams.” — Goldin, Venneri and Noor [1]

The United States (U.S.) governments National Aeronautics and Space Administration's (NASA) programs and missions are widely dispersed among various NASA Centers and contractors. Historically this has implied a geographic centralization at a NASA Center of skills required for a mission to be completed. Today's NASA activities, almost by definition, are highly distributed in nature. Currently they are distributed across the United States. However more and more they are distributed around the world and may soon pass those limits with the development of the Deep Space Internet. This distribution is coupled with a trend towards the development of focused expertise in narrower areas with that expertise residing in fewer people who are distributed

across a wider area. The application of one of these experts to a problem within their area of expertise brings dramatic results in a very short period of time. The challenge is to effectively share these abilities across great distances and create highly distributed virtual teams containing the necessary expertise to support NASA's missions.

As we enter the 21st century, the need for more productive distributed engineering environments will greatly change the way we engineer systems. No longer do we have the option of slow design cycles and separate component engineering for aerospace systems. The mantra “faster, cheaper, better” is required by the geographic distribution of talent, the U.S. funding realities, the high quality, and rapid mission requirements of today programs. As Goldin, *et. al* have observed, we need intelligent synthesis environments that allow engineering tools to be freely applied to the *virtual* life-cycle and analysis of NASA products. Appropriately, then, NASA has embarked upon a new initiative called the Intelligent Synthesis Environment (ISE) [2]. This initiative will prototype the applications, environments, and tools, for distributed design, operations, and support of NASA missions and programs. An outgrowth of the ISE initiative is the IsoWAN activity.

Virtual life-cycle implies the capability to assess and analyze the impacts and variables of design and operational decisions without the development (or with the very limited development) of hardware prototypes. This distributed, collaborative, simulation of a life-cycle of an aerospace is within the technical capability of the existing analysis tools except for the fact that these tools tend to be hand-crafted, finely tuned, built to operate in a unique phase of the life-cycle, and to operate within well formulated computational environments. An intelligent synthesis environment requires something else.

Traditionally, NASA has built engineering analysis and design tools for specific computing architectures and purposes. Any distribution, sharing, or composition of such elements was explicitly part of their design and laboriously constructed to the particular architecture and topology available. NASA wishes to create a general capability for distributed applications. Such applications demand application services interfaces (ASIs) that are programmable interactions between applications and the

wide area services. The state-of-the-art in providing such distributed services exists in the World-Wide Web, where the user has the illusion that a remote system or functionality is local to an application's environment, and can be accessed and utilized like any other local service. While this works well for web-surfing and shopping, more complex applications require invoking a greater variety of services. Examples of some of NASA's desired distributed services [3] include a Public Key Infrastructure (PKI), virtual private networks (VPNs), directory (i.e. X.500/LDAP) services, product data management services, distributed file management (i.e. DFS), database, and audio/videoconferencing/whiteboard (i.e. H.323 / T.120) services.

Distributed application services also allow applications to act in a location independent manner over NASA's distributed wide-area network infrastructures, but is only a first step towards a NASA-wide virtual life-cycle environment. Users and contributors to a distributed application of shared data and tools will demand appropriate policies on issues such as security, reliability, data ownership, quality of service, and management and monitoring of processes. These will be discussed in a later paper. This paper discusses architectural coordination and modifications to NASA's wide-area and local-area networks and service infrastructures to enable and support secured, wide-area application service delivery in support of NASA missions.

Distributed service frameworks, like the IsoWAN, evolve up from the basic network infrastructures which have traditionally offered only data transfer to customers. Advanced networking capabilities, such as VPNs, Quality of Service (QoS), and resource scheduling, advertisement, and routing via directory control are changing the nature of how wide-area networks can be utilized. These technologies are enabling networks to understand and host distributed services based on traffic flow types or application services used, coordinate the delivery of services and data to the customer, and dynamically couple users or user groups to applications. In considering the network from traffic flow, application services, or user/user group perspectives, the traditional physical boundaries of networks – local-area through wide-area; intranets and extranets - are superseded by the end-to-end requirements of the users, their applications, the supporting application services, and the resulting traffic flows. The nature of these end-to-end customer and service requirements, as well as the enabling characteristics of emerging virtual network technologies, prompt us to consider new network architectures to support application services delivery and interfaces.

Existing networks are architected to provide an infrastructure for basic data delivery, including packet forwarding, routing and peering, domain naming service (DNS), and other network-layer core capabilities (e.g., policies, web caching, multicast, virtual networking). The

concept of providing higher-level application services to customers of the network is recognized as important, but as yet there is little work in that area. In addition, there is no support in existing networks for emerging environments such as secure group collaboration, distance-independent computing, grid computing [4], collaborative engineering, or intelligent synthesis environments. The nature of environments such as these is that the requirements that drive the architecture are primarily end-to-end, and will require the network and application services systems to understand and provision the services based on those requirements, as well as coordinate service delivery between networks and to the customer.

A logical starting place to focus on service delivery is at network interconnect points, such as Network Access Points (NAPs) and Isolation LANs (IsoLANs) for NASA Centers. Network interconnect points act as distribution networks, where access and backbone networks peer with each other, and where secure distributed services and application systems can be located to optimize application traffic flows. However, the architectures of current network interconnect points are rudimentary, providing primarily basic routing and peering functions between access and backbone networks. Current network architectures are not well suited to support services above the network layer, nor to support dynamic couplings of users and applications to these services.

2: IsoWAN Concept

The next evolutionary step in supporting wide-area applications and services delivery to customers is a network framework that provides for collocation of applications and services at distinct sites in the network, an interconnection between these sites that is performance optimized for these applications and, value-added distributed services and interfaces for applications. The IsoWAN describes a set of isolated wide area network nodes, and is an advanced, secured interconnect framework for services that will make applications easier to access in distributed environments. The main functions of an IsoWAN are localization of application services, coordinated delivery of applications and associated data to the customer, and support collaborative application development for customers. The IsoWAN is being implemented, in part, to support NASA's Information Technology Base Program [5], the NASA Chief Information Officers, and the new Intelligent Synthesis Environment (ISE) Initiative's distributed applications. Three NASA Centers have joined together in a pilot implementation of the IsoWAN. These Centers are the Ames Research Center (ARC), the Jet Propulsion Laboratory (JPL), and Marshall Space Flight Center (MSFC) has been built. This network will be expanded

to all 10 NASA Centers to form the full-fledged IsoWAN for NASA.

The basic concept of the IsoWAN supports distributed application services by the creation of local Center specific “nodes” that are interconnected to form unified, yet distributed access points. Each of the IsoWAN nodes is a local proxy point for all of the applications and services proposed. They are NOT intended to act as gateway points for wide area networks (WANs), as those already exist. They are placed in-between the WAN and the NASA Center local isoLANs to provide localized access to the applications and services. Figure 1 shows the basic concept of the IsoWAN node and its connection to the other IsoWAN nodes at various NASA Centers across the U.S.

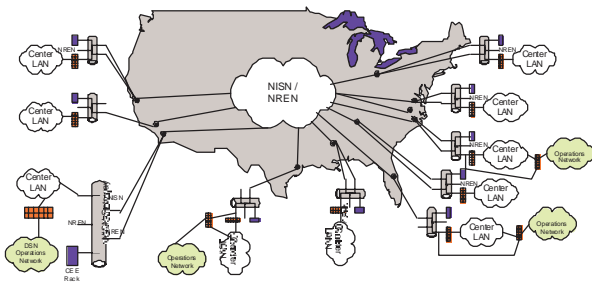


Figure 1 – IsoWAN Network Topology

IsoWAN services hosted at the nodes (above the network layer) are defined here as value-added support of distributed applications for users. This includes collaborative services, distributed management of users or user groups, and coordinated secure delivery of multiple applications along with their associated data to a customer. The local NASA access to these services and applications will be through the IsoWAN node hosted at that site. All access to NASA IsoWAN services and applications will be provided via an appropriate and document ASI that will be identical, regardless of which NASA Center the user is accessing from. This is similar in concept to grid computing [4] except that instead of a focus on distributed supercomputing and queue management, the focus is on more user-centric and considers collaboration services and security.

An application service consists of value-added support for the user's and applications. This includes supporting communication between the application server(s), and the user, consisting of:

- Coordinating secure delivery of the application and associated data across all networks between customers and/or between the customer and server
- User directories to control preferences, application access, and group-based user policies.
- File systems that are scalable across and between NASA Centers, such as AFS or DFS
- Application and user data movement and management

(transmission across backbone, mirroring, staging, caching, migration, backups)

- Provisioning of network bandwidth for application and data delivery
- Monitoring the status of the application session, and providing billing and accounting information
- Scheduling, prioritization, and preemption of application sessions
- Prioritize and synchronize applications with real-time multimedia user collaboration sessions.
- Scheduling the delivery of non real-time multimedia data.

Services for multiple concurrent applications will be exposed via a software set of application services interfaces. These ASIs will enable the distributed above listed services without having to tie to a specific technological implementation. The intent is to provide an environment where multiple distributed applications are supported by the IsoWAN framework and multiple applications are experienced concurrently by a variety of users throughout NASA.

3: Application Example

One of the first uses of the IsoWAN framework is in support of the across the NASA Collaborative Engineering Environment (CEE) rooms developed by the ISE initiative [6]. The CEE rooms are a production test area for new tools and technologies. IsoWAN is one of the first frameworks to support the CEE rooms, and is expected to be the foundation upon which current and next generation CEE capabilities and services are built.

The service levels essential to this level of collaboration are a very large step from those currently needed to support e-commerce. The stakes are several orders of magnitude higher. The basic product being managed is still information, but the value associated with what is being supported is measured in lives, multi-billion dollar national resources, and corporate profits. Classically, those who work in this arena do so in secure research centers or, when distributed teams are utilized, over highly secure and dedicated communications links. However, while this point to point model supports a single, compartmentalized, distributed team, it falls short of supporting a topology where each expert is a part of multiple teams that are distributed across the country or across the world.

As a first step in addressing the needs of these teams, Collaborative Engineering Centers have been installed at all the NASA Centers and the baseline services for a collaborative infrastructure have been defined. In addition, the necessary methodologies for the effective application of these tools, as well as support for the necessary cultural change, are being addressed by ISE in the design phase.

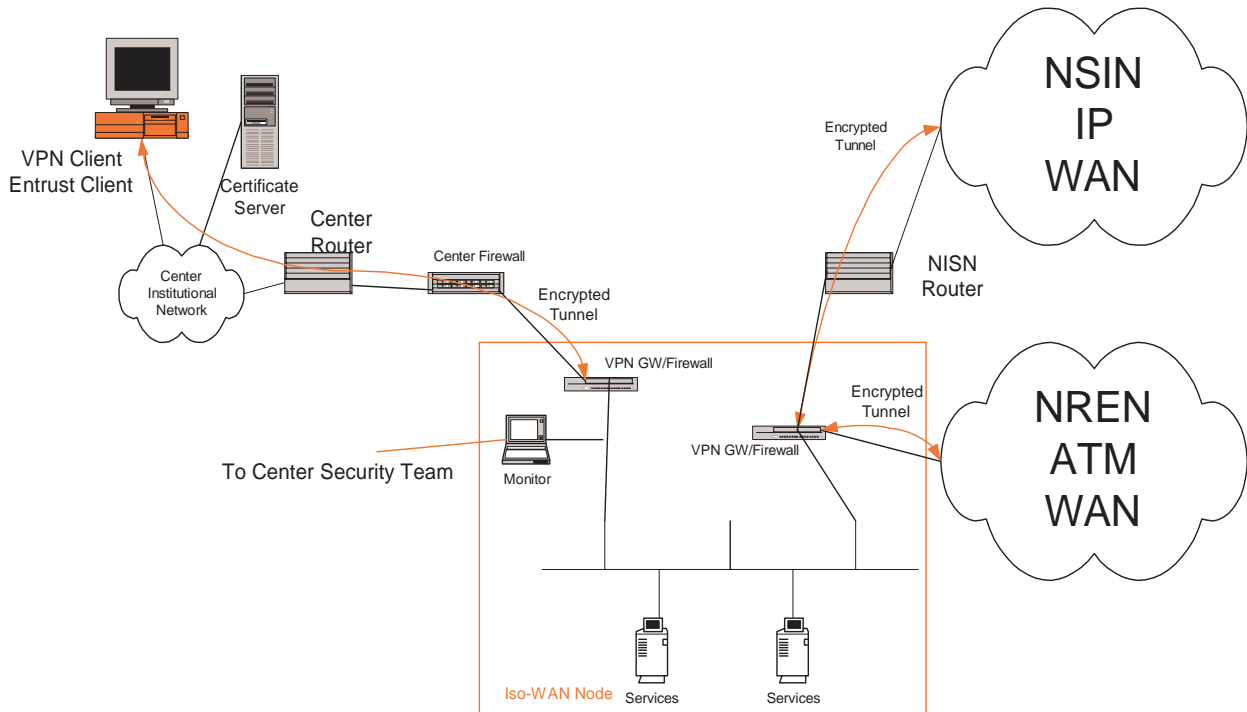


Figure 2 – Two Secure Network segments into and out of the IsoWAN node

The basic capabilities provided by the initial IsoWAN nodes: security, conferencing, shared storage and connectivity; create a collaborative infrastructure that will begin to meet some basic distributed team needs. Initial utilization will include "classical" collaborative activities such as audio, video and data conferencing. First use will be to connect research, development, implementation and operations teams from multiple NASA Centers as well as the provider and customer community.

The goal is to provide better and more supportable life cycle designs, quicker and for a lower cost. A major focus, and probably the most important early benefit, is the reduction in travel related costs, specifically in the area of "lost time". In today's smaller and smaller organizations, there is very little capability duplication. An expert can either be at the control center, supporting the current mission; at the launch site, supporting the next mission; or at the plant, supporting the next project. The expert cannot be in all three locations at the same time. A week spent between the various sites will consist of the majority of the week spent traveling, and a minority of the time spent adding value to the process. The ability to quickly and securely connect the expert to all of these locations, and to provide for a quality of collaboration sufficient to replace collocation, can easily double the effectiveness of the expert through simply doubling the amount of time spent adding value to the efforts. Other increases are also available due to reduced fatigue, the fact the expert has full access to their normal resources, and the ability to easily add non-expert supporting personnel to the effort.

Initial application efforts at Kennedy Space Center (KSC) have taken advantage of the existing NASA CEE rooms around the NASA. Follow on efforts are centered on the recently defined mini-CEC. There are currently plans to deploy several of these at KSC in the year 2000 to extend the collaborative reach from the manufacturing factory to the launch site and pick up all of the locations in-between. IsoWAN, and application services are what will make this possible. By treating and managing IsoWAN supported applications as a shared collaborative infrastructure, as opposed to a project related activity, NASA is going to be able to significantly leverage a rather small investment into a national capability.

Follow on use is intended to provide a shared, multi-user, secure modeling environment that will allow engineering and analysis to occur concurrently across the relevant community with a managed, shared and maintained set of information. This effort is expected to start with facility and infrastructure models and rapidly progress to launch vehicle, spacecraft, launch infrastructure and space environment models that are globally validated, utilized and shared. Initial discussions with other potential users indicate plans for similar first use efforts up to and including partner arraignments to develop the initial information base.

An significant benefit provide by ISE as well as IsoWAN will be the overall cost reduction associated with this type of work due to shared infrastructure, services, models and methods. The reduced costs will allow application to the smaller design and development activities that constitute a significant portion of NASA's

development efforts.

4: IsoWAN Pilot Network

The development of an enterprise-wide collaborative information infrastructure such as the IsoWAN is a complicated process. It requires buy-in and coordination from geographically distributed NASA Centers that have their own policy and procedures on how restricted engineering/scientific/business information can be shared. Some of the advocacy that IsoWAN development team went through are as follows:

1. Gathered user application requirements.
2. Developed straw man architecture and basic services based on user requirements.
3. Obtained enterprise networking organization's support for future deployment.
4. Obtained all NASA Center Chief Information Officers' support.
5. Obtained NASA security team's support.
6. Obtained each Center's institutional networking team's support.
7. Developed relationships and obtain inputs from NASA's industrial and academia partners.
8. Held two IsoWAN workshops to revise IsoWAN architecture and services.

The final IsoWAN implementation is based on the inputs from all key participants. IsoWAN development has been a phased implementation, with the first phase a

pilot network between three NASA Centers. An initial IsoWAN node will be created in the form of a rack that will be placed in the wide area network gateway in every NASA Center. NASA's operational wide area network services are provided by the NASA Internet Services Network (NISN). NISN is an outsourced contractual service for raw bandwidth and managed basic IP services such as routing and Domain Name Services.

In the initial pilot the IsoWAN node, logically, sits between the Center's firewall and the WAN network access point. The idea is to have a physically secured area just outside of each Center's firewall so that distributed services can be placed in this secured area. Access to the services in this secured area is strictly restricted to authenticated users through network level security, i.e. an encrypted network tunnel. As shown in Figure 2, encrypted tunnels are built between the IsoWAN nodes through NASA's wide area networks. Another set of encrypted tunnels is built between the IsoWAN node through the Center's firewall to the users in each Center. Users have to be authenticated using the NASA's X.509 certificate system before the encrypted tunnel can be built. The tunnel can only go from inside the Center's firewall to the specific IsoWAN node that is trusted by the Center. The data inside the IsoWAN node is not encrypted. A special physically/logically secured link from inside the IsoWAN is provided to each Center's security team to provide an extended monitoring/control point before the data is allowed to go through the encrypted network link.

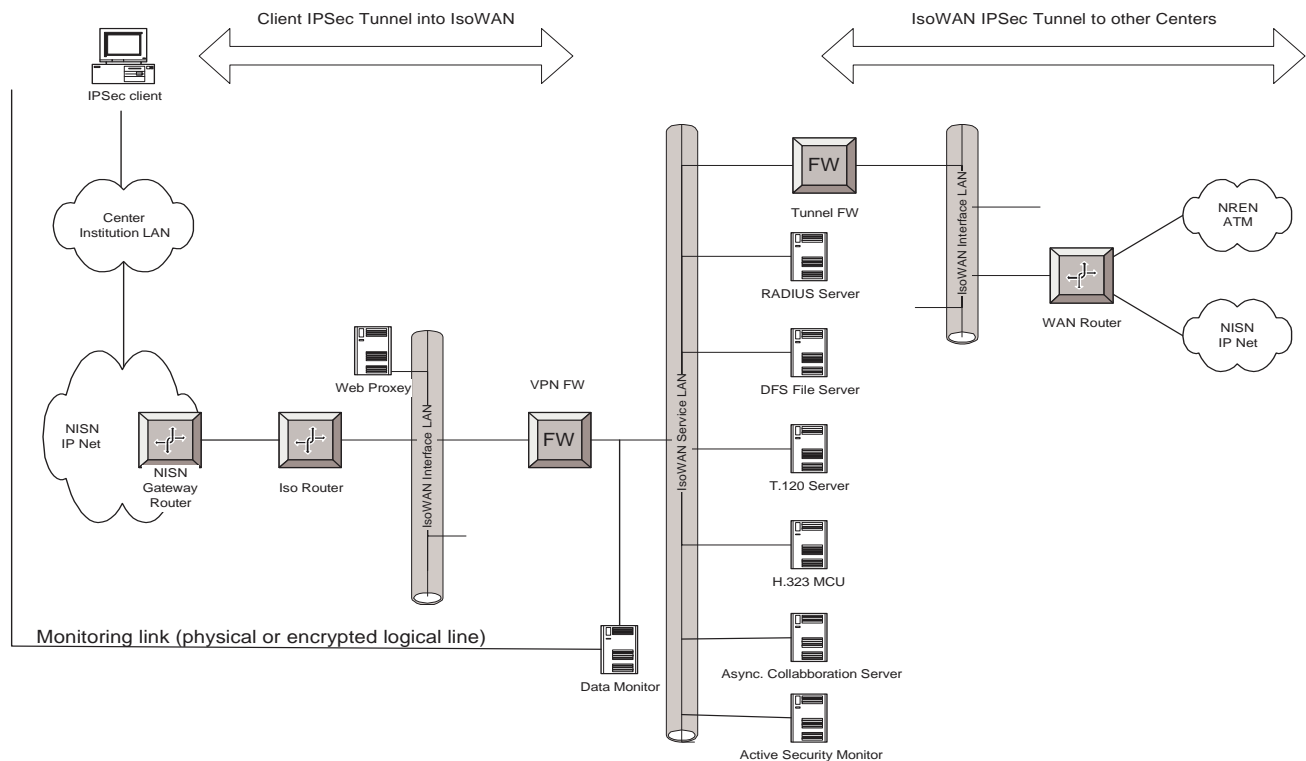


Figure 3 – IsoWAN Node Infrastructure

This two segment, encrypted network tunnel architecture permits the desired virtual localization of application services. It also provide the advantage that, since data is monitored before encryption, advanced protocols such as T.120 [7], H.323[8], and other application specific protocols can go through the Center's firewall using the encrypted tunnel. This function is critical for ISE and other NASA programs, which will develop a number of distributed collaborative applications each of which may use several different high-end protocols.

Figure 3 shows the block diagram of the IsoWAN node in the three nodes pilot network between ARC, JPL, and MSFC. IPsec technology [9] is being evaluated to encrypt traffic between IsoWAN nodes, and between IsoWAN node and users. Because the main objective is to identify the networking, performance, and security issues, only limited application services and ASIs are provided in the pilot network.

The services that will be supported in IsoWAN will evolve with the user application requirements. Some have been previously mentioned in this paper. The initial services and ASIs during the IsoWAN pilot are provided for user access through a user customizable Web portal that proxies the services from the IsoWAN node. Users can then access the services on any machine with a Web browser. Some example use scenarios are:

A group of distributed users can securely collaborate using T.120 application sharing to view the same CAD drawing. They can also use the H.323 to provide multipoint audio/video conferencing bridging.

A distributed design team can use the document management service to control revisions of documents. They can also use the DFS to provide distributed data sharing. The automatically replicated DFS binary storage can assure the team will use the same versions of data and application software as well as limit wide area data transfers to only that needed to preserve the information. The product data management services can provide access to different design databases for knowledge reuse and manage product development from modeling to production.

For a team doing distributed product-modeling simulation, the DFS can also help by providing fast, localized data access. Because NASA projects often last several decades, DFS with a consistent data path and name structure will help make data available to operations teams to allow design and operational simulations to be performed decades later. The directory-based networking service will automatically give higher priority to users and applications to reduce the latency for simulation data.

5: IsoWAN Status and Plans

The three-node IsoWAN pilot network provides us with a platform to test the network performance, security issues, and services performance. A number of potential problems have been identified through this pilot network. For example, VPN technology using address translation has created IP address security problems for DFS. Providing IP addresses from within the IsoWAN node also present problem with scalability. A routed IPsec protocol was used instead of VPN, which resolved the above problems. Another example is that the IsoWAN services are distributed services connected over wide area network so the network latency time can create performance problem for some of the services. The IsoWAN team is carefully optimizing the network performance of each service.

The current plan is to complete the NASA IsoWAN development at the beginning of 2001. The system should go operational by mid to end of 2001.

6: Conclusion

This paper presented NASA's plan for a science and engineering information and services framework called the IsoWAN. This collaborative information infrastructure is a virtually collocated service based framework. It will enable distributed team collaboration and support *virtual* life-cycle analysis of NASA products and missions.

7: References

- [1] Goldin D. S., Venneri S. L., Noor A. K., *Beyond Incremental Change*. IEEE Computer 1998; 31(10): 31-39.
- [2] ISE Initiative Transition Office, *NASA's Intelligent Synthesis Environment Initiative, Level 1 Plan*, - Draft, NASA Langley Research Center, 1999.
- [3] CEE Infrastructure Team, *CEE Infrastructure Functional Requirement Document*, Draft 1.1, NASA, 1999.
- [4] Ian Foster and Carl Kesselman, *The Grid, Blueprint for a New Computing Infrastructure*, Morgan, Kaufmann, 1999.
- [5] IT Base Program Office, *NASA's Information Technology R&T Base Program, Level 1 Program Plan*, NASA, December 1999.
- [6] Launch Space Magazine, *KSC and NASA's Intelligent Synthesis Environment*, March/April issue, page 50, 1999
- [7] ITU-T, *T-120 Standard*, International Telecommunications Union (ITU), 1998
- [8] ITU Study Group 16, *H.323 Standard Version 2*, International Telecommunications Union (ITU), 1998.
- [9] William Stallings, *Cryptography and Network Security*, Prentice Hall, 1999.